Guiding Technology Deployment Decisions using a Quantitative Requirements Analysis Technique

Martin S. Feather*, Kenneth A. Hicks*, Ryan M. Mackey*, Serdar Uckun**
Jet Propulsion Laboratory, California Institute of Technology
***NASA Ames Research Center*
{Martin.S.Feather, Kenneth.A.Hicks, Ryan.M.Mackey}@jpl.nasa.gov, uckun@mail.arc.nasa.gov

Abstract

Successful and convincing operation of a prototype, deployed in a real setting, is a key step in advancement of many a new technology from research laboratory to real-world use. Often, however, such a deployment must be interjected into a pre-existing context of ongoing activities, established designs and standard practices. That context can pose a number of obstacles, which if unaddressed can preclude success. Careful selection of what demonstration opportunities to pursue, and determination of how best to pursue them, are therefore crucial.

A study was conducted to select and plan for deployment of prototypes of integrated system health management (ISHM) software on NASA spacecraft. The study itself utilized our seasoned technology maturation assessment process, based on a quantitative requirements analysis technique. However, this process is typically applied to scrutinize a single technology application at once. In this case there were a number of candidate deployment opportunities. Since it would have been tedious and time-consuming to consider each of them one-by-one, we adapted our assessment process to accommodate their simultaneous consideration. We relate our experience in doing this – the shortcuts we took, the similarities we exploited, and the workarounds we adopted to complete this study in a timely yet effective manner.

1. Introduction

The focus of this study was a range of emerging technologies of interest to NASA missions. The next step in their maturation towards full-scale adoption by a space mission would be deployment and operation of a prototype flying as an adjunct on an actual mission. However, many obstacles face such deployments, and selecting the right deployment opportunities, and planning for them, must be done with care to maximize their chances of success.

This paper describes how we adapted our technology assessment process to guide such selection and planning. Section 2 summarizes the range of technologies, Section 3 summarizes our technology assessment process, Section 4 describes why and how we adapted that process to match the needs of this study, and Section 5 concludes the paper.

2. Integrated System Health Management (ISHM) technologies

Integrated System Health Management (ISHM) technologies are intended to augment a system so as to prevent, mitigate, and/or recover from faults occurring during the system's operation. NASA has an interest in such technologies as a cost-effective way to increase the reliability of its spacecraft.

2.1 ISHM background

The following is a highly condensed version of the summary of ISHM's evolution described in [1]:

NASA's interest in this area began in the early 1990s, in the "Vehicle Health Monitoring (VHM)" work that focused on selection and use of sensors and software to monitor the health of space vehicles. "System Health Management" became the preferred moniker when this was extended to complex human-machine systems, not just the vehicles themselves, and to consider what actions to take based on the parameters monitored. Similar work taking place under the auspices of the Department of Defense was referred to as "Integrated Diagnostics," where the focus was on operational maintenance issues (usually in an aircraft environment). When detecting and responding to faults requires "integrated" diagnostics that look at many aspects of the system in question, the terminology "Integrated System Health Management" (ISHM) applies. Further information on ISHM is in papers from a forum devoted to the topic, available from
http://ti.arc.nasa.gov/projects/ishem/papers_pres.php

2.2 ISHM for upcoming NASA missions

ISHM technologies have had success in supporting the operational maintenance phase of, for example, commercial aircraft and trains. These are arenas where there is availability of historical information on the operating characteristics of many instances of the same system. NASA's upcoming missions will not have this wealth of heritage data. Instead, they will be spacecraft with few prior equivalents. ISHM technologies offer the prospect of improving the reliability of such systems, but lack the track record to establish this. A prototype of ISHM technologies, flying as an adjunct on an actual mission, is needed to mature them towards acceptability by space missions. In the Technology Readiness Level scale customized to spacecraft information technologies [2], success of such would raise them to TRL 7 – "*Information technology validated in space. Adequate documentation prepared for transfer from developers to full operations engineering process team.*" However, such a prototype deployment promises little in the way of benefit to its host mission, so must bear the brunt of accommodating to that mission, during both development and operation. It falls upon the proponents of the prototype to resolve the many obstacles that arise.

The purpose of this study was to select the most promising from among a variety of opportunities for ISHM prototype deployment, including planning (as could be afforded) of extra steps to take to help overcome adoption obstacles. There seemed little in the way of standard methodologies for such planning and selection, so the team decided to adapt an existing technology assessment process for this purpose.

3. The Technology Infusion Maturity Assessment (TIMA) process

"Technology Infusion Maturity Assessment" (TIMA) is a process developed at JPL and used for several years for assessing and planning the maturation of (predominantly spacecraft) technologies [3].

3.1 TIMA background

The TIMA process was designed to improve the success rate for maturing novel technologies to the point where they become acceptable for use on space missions. It helps determine the requirements that the novel technology will need to fulfill, and determine the way to advance and demonstrate sufficient maturity of the technology with respect to those requirements.

The TIMA process is conducted in facilitated group sessions, during which information is elicited from stakeholders and combined to yield a model that supports the determinations outlined above. At the core of the TIMA process is a quantitative requirements analysis technique initially conceived as an aid to quality assurance planning [4]. In this technique, a model is comprised of instances of three kinds of concepts: *Objectives* – the requirements or goals (functional and quality – a.k.a. "non-functional") of the project, system or technology, and its development, *Risks* – what could occur to impede the attainment of Objectives, and *Mitigations* – what could be done to reduce Risks. Quantitative relationships connect these concepts' instances. Risks are connected to the Objectives they threaten (indicating by how much the occurrence of a Risk would detract from attainment of an Objective), and Mitigations are connected to the Risks they reduce (indicating by how much the application of a Mitigation would reduce a Risk). The constructed model allows for investigation of the costs (of the selected Mitigations) and benefits (calculated in terms of Objectives' attainment) of alternate selections of Mitigations. By viewing obstacles to technology infusion as Risks (potential future events which, should they occur, will impede infusion), the methodology becomes applicable to assessing and planning the infusion of technologies. Hence, it seemed a good process to try to adapt to the needs of the ISHM study.

3.2 TIMA for ISHM deployment decisions

In seeking to apply the TIMA process to the challenge of making ISHM technology decisions, we faced a significantly different situation to the usual TIMA study, applied to a technology aimed at a *single* mission or a homogeneous class of missions. By way of contrast, the ISHM deployment study needed to simultaneously consider *multiple* significantly different deployment opportunities.

Our initial list of such identified 21 such opportunities. It would have been far too onerous to separately apply the TIMA process to each. While there might be some savings to be had from not having to explain the TIMA process to a new group of people each time, total effort would nevertheless be substantial (a typical TIMA study of a single technology takes on the order of four half-day sessions, each session requiring participation of nearly all the key stakeholders). Thus we needed to streamline the TIMA process – how we did so is described next.

4. The experience of adapting TIMA to the ISHM deployment study

The study assembled a team of relevant ISHM subject area experts drawn from JPL and NASA Ames.

They followed the TIMA process, adapted (primarily to streamline it) as described in this section. Over a period of two months the team convened in several meetings (each of several hours duration), interspersed with short teleconferences and email exchanges. Post-session analyses, and generation of documentation, were performed in the following month.

4.1 Narrowing the field

The initial scope of the study was a list of 21 mission opportunities identified by the team as potential targets for deployment of ISHM prototypes. The target missions varied substantially in terms of their maturity levels, customers, and overall budget. Early on the team realized that it would be infeasible to consider all 21 in depth, so the first step was to find a way to narrow the field. This was done by retaining only those opportunities for which one or more individuals gave an indication of eagerness to champion its pursuit. The was some risk to this step – among the opportunities discarded from consideration there might have been one that, had it been studied, would have emerged as a preferred opportunity; it is plausible to imagine that a champion for pursuit of that opportunity could then have been recruited.

This narrowing resulted in retention of 7 out of the original 21 deployment opportunities, a number that the team felt could feasibly be scrutinized in detail using a suitably streamlined TIMA process. These 7 are outlined below (but for sensitivity reasons we do not identify the specific missions):

1. Data from the flight of a launch vehicle would be analyzed using ISHM technology after the flight was over.
2. Data during a launch vehicle's preflight checkout period would be analyzed in real-time using ISHM technology.
3. Data would be monitored and analyzed during flight of a launch vehicle by ISHM algorithms executing on a dedicated on-board computer, with outputs written to a data recorder for post-mission recovery.
4. ISHM in a ground-based fault prediction console would help with the ground testing of a spacecraft through Assembly Test and Launch Operations (ATLO), and during flight to monitor sensor readings so as to proactively identify problems.
5. Establishment of an airborne testing facility capable of recreating Mars levels of gravity, temperature and atmospheric pressure.
6. An on-board active ISHM experiment on an earth-orbiting spacecraft – exercising of ISHM (by deliberately commanding failure modes, etc) on the spacecraft once it has completed its primary mission.
7. Use of on-board active ISHM for fault protection during an earth-orbiting spacecraft's primary mission, and, once the primary mission is complete, for more ambitious uses of ISHM including experiments on ISHM directing autonomous spacecraft control while faults are deliberately injected.

Note that while ISHM technology is the recurring common element (except perhaps for #5, which emerged as somewhat of an oddity), in other respects these deployment opportunities differed considerably.

4.2 Focus on obstacles, ignoring merits

The decision was made to streamline the TIMA study to focus on only the *obstacles* to ISHM prototype deployments, not the merits that (successful) execution of those prototypes would illustrate.

Normally a TIMA study would elicit the requirements that adoption of a new technology would help fulfill. This typically helps in several ways – ensuring the technology is well-matched to the mission needs, prioritizing the requirements that technology is to fulfill, and determining the best allocation of requirements between technology and its environment. In this ISHM study however, most of the deployment opportunities were intended to allow an ISHM prototype to operate unobtrusively alongside the actual mission, without the mission relying on that prototype (other than relying on it to remain unobtrusive!). Thus there was much less of a motivation to consider the requirements that the prototype would fulfill.

The shortcoming of ignoring merits is that the TIMA study would yield an understanding of only the *obstacles* to each of the deployment opportunities. Due to the significantly differing nature of those opportunities, successful operation of the ISHM prototypes would yield significantly different levels of understanding of the *merits* of those technologies. Even a deployment that failed to achieve all of its objectives would likely yield some increase in understanding useful for advancement of ISHM. Nevertheless, the team felt that they could separately gauge the relative benefits of each of the deployment opportunities. Their most pressing concern was to focus on obstacles, where the greater uncertainty lay (including uncertainty about how best to plan to overcome them).

4.3 Assume a baseline of "standard practice"

The TIMA study considered the obstacles that each of the deployment opportunities would face, and identified actions that could be taken to help overcome those obstacles.

In making these assessments, an assumption was made of a baseline of "standard practice" for prototype deployment – i.e., the course of action that would normally be pursued in deploying a prototype of a novel technology (many of the participants were familiar with such, having long been involved in prototypes and their deployment). Normally a TIMA study would scrutinize the details of such "standard practice", and explicitly assess its effectiveness at overcoming obstacles. The usual purposes for doing this are twofold – to reconsider the efficacy of standard V&V practices as applied to a novel technology (the novelty may mean that some practices' efficacy differs significantly from what one would expect), and to capture the purpose and extent of reliance on those practices (so that if there is the need to make some cutbacks on what can be performed, the increase in risk can be ascertained and minimized).

By making this baseline assumption, the overall contribution of standard practice would be retained, built-in to the team's assessments of obstacles' magnitudes, yet time would be saved by avoiding the need to go into the details of standard practices. The risks of shortcutting this step were felt to be of lesser significance to this ISHM deployment study than would normally be the case. The team's experience with prototypes gave them confidence to make such assessments. Meanwhile, the desire to entertain cutbacks of standard practices was of secondary consideration when compared with the more pressing concern, namely identifying (and addressing) novel obstacles that arise in the case of the various deployment opportunities.

4.4 Representing generic and specific information

As is typical for a TIMA study, the bulk of the effort went into building the TIMA model of the situation. This involved:

- identifying obstacles that each deployment opportunity would face,
- assessing the magnitudes of those obstacles (in the TIMA approach, these assessments have a coarse quantitative nature – each obstacle is assessed against each deployment opportunity to capture an estimate of the expected magnitude of the impediment to success it would pose),
- identifying the actions that could be taken to help overcome those obstacles,
- assessing the effectiveness of those actions (in the TIMA approach, these assessments also have a coarse quantitative nature – each action is assessed against each obstacle to capture an estimate of the expected reduction of the obstacle that performance of that action would achieve), and
- assessing the costs of each of those actions.

This model building was done for all seven of the deployment opportunities at once, starting with the obstacles and then moving on to consider the actions. Having identified an obstacle to one deployment opportunity, it was straightforward and speedy to consider whether it would impede each of the deployment opportunities, and if so assess by how much. Each deployment opportunity was represented as an "Objective" within the TIMA model, and each obstacle as a "Risk"; the magnitude of each obstacle against each deployment opportunity was represented in the TIMA model's relationship between the two.

The analogous step for actions was to represent each as a "Mitigation" in the TIMA model. However, this approach proved problematic when an action's effectiveness at overcoming an obstacle differs depending on which deployment opportunity was under consideration. This distinction violates an assumption of the underling TIMA modeling framework, namely that a Mitigation has the *same* effect on a Risk (e.g., reduces it by 90%) for each and every one of the Objectives that Risk threatens. Thus the way we had intended to model ISHM deployment (opportunities as Objectives, obstacles as Risks, and actions as Mitigations) would fall prey to this limitation. Our workaround for this problem was to make a separate copy of each of the obstacles, one per deployment opportunity. For a given obstacle copy, its magnitude would be scored as non-zero against only the one deployment opportunity to which it applied. Meanwhile, an action's effect on that obstacle copy could have one value, while that same action's effect on a different copy (applying to a different deployment opportunity) could have a different value.

The way we gathered this information was to ask first whether the action had the same obstacle-reducing-effect regardless of deployment opportunity. It the answer was "yes", we would make the assessment of that effect, and move on to the next action. Only if the answer was "no" would we decompose the estimation depending on the deployment opportunity.

The net result of this was the ability to capture the distinctions that the team felt existed, but at the cost of a verbose underlying model (resulting in disadvantages of being clumsy to scrutinize and maintain).

4.5 Modeling cost distinctions

Another area of mismatch between the TIMA modeling framework and the nature of the ISHM study arose when it came to costing the actions. We realized there were three forms of costs:

- Actions that only need to be paid for once, and thereafter yield their benefit to every deployment opportunity without any additional cost.
- Actions that have to be paid for on a deployment-by-deployment basis, in turn subdividing into:
    - Actions whose cost is the same per deployment
    - Actions whose cost differs from one deployment to another

Again, we adopted a workaround based on copying of information when necessary. Further details are omitted here in the interests of brevity.

At the end of this phase, the TIMA model of the deployment obstacles and potential actions to overcome them had been completed.

4.6 Decision making and off-line analyses

In most TIMA applications, model building is completed by the end of the penultimate session. Computationally intensive analyses are then run off-line, after which the team is re-convened to make decisions based on scrutiny of the analysis results, and with access to the TIMA model from which they are derived (e.g., they can try "what if" scenarios, switching one action, say, for another). In this study there was insufficient time left to reconvene the whole team yet again, so instead the team made some decisions using just the model (without the analysis results), and relied upon post-session analyses to confirm the validity of their decisions. The approach worked well for this study, perhaps because the simplification of the overall model to omit a TIMA "Requirements" (section 4.2) made it more feasible to manually arrive at wise decisions. The way the team pursued this was as follows:

For each opportunity, the team used the model to help them identify a "prudent" set of actions to follow. This identification was deliberately frugal, aiming to limit total cost to amounts commensurate with their intuitions for what expenditures each of the deployment opportunities would warrant. The primary use of the TIMA model was to help them identify relatively low-cost ways to overcome major obstacles. They did this identification in the following steps:

- The team considered obstacles in decreasing order of their expected magnitude, stopping when all the remaining obstacles were of low magnitude (which was deemed to be below approximately the 10% level of obstruction).
- For each obstacle, the team selected actions that were effective (i.e., had a significant effect at overcoming that obstacle) and cost no more than approximately 100$K. (Note that an action may simultaneously help overcome multiple obstacles, a phenomenon that the software supporting the TIMA process made visible to the experts.)
- After selecting all such low-cost actions, the team added more expensive actions only if they were highly effective at addressing critical obstacles that the team felt could not be left outstanding.

After these sessions were completed, off-line analyses were run to explore the cost-benefit tradespace, using the simulated annealing optimizer built-in to TIMA's software. These analyses confirmed that all but one of the "prudent" selections were close to optimal, that is, for the same amount of money, no significantly superior selection of actions (superior in the sense of leading to greater reduction of obstacles) could be found. There was one case of a non-trivial improvement identified by the off-line analysis, which, with the concurrence of the study leads, replaced the team-identified "prudent" selection.

These off-line analyses were also used to confirm that the "prudent" selections were at reasonable locations on their respective Pareto frontiers – i.e., the team had not overlooked an opportunity to achieve a lot more reduction of obstacles with only a little increase in expenditure. In one case a prominent "step" further along the frontier was revealed – for essentially double the cost of the team-identified "prudent" selection, the remaining level of obstacles could be halved. This was called out in the final report as an alternative.

For reporting purposes, the TIMA software was adjusted to generate a single cost-benefit chart of the Pareto frontiers for all seven of the deployment opportunities (shown on the next page). Its vertical axis is a measure of total obstacle level, so lower is better; its horizontal axis is cost, so to the left is better. The wiggly lines indicate the Pareto frontiers (optimal selections of actions to most cost-effectively reduce obstacles). The circles indicate locations of the "prudent" selections.

Note that opportunity #5 is by far the best in terms of low cost to minimize obstacles – as mentioned earlier, it was somewhat of an oddity, representing a facility rather than an ISHM technology per se. At the other extreme, it is clear that opportunity #3 is the worst from a cost/obstacles perspective. However, recall that we deliberately ignored the relative merits of each of these opportunities.

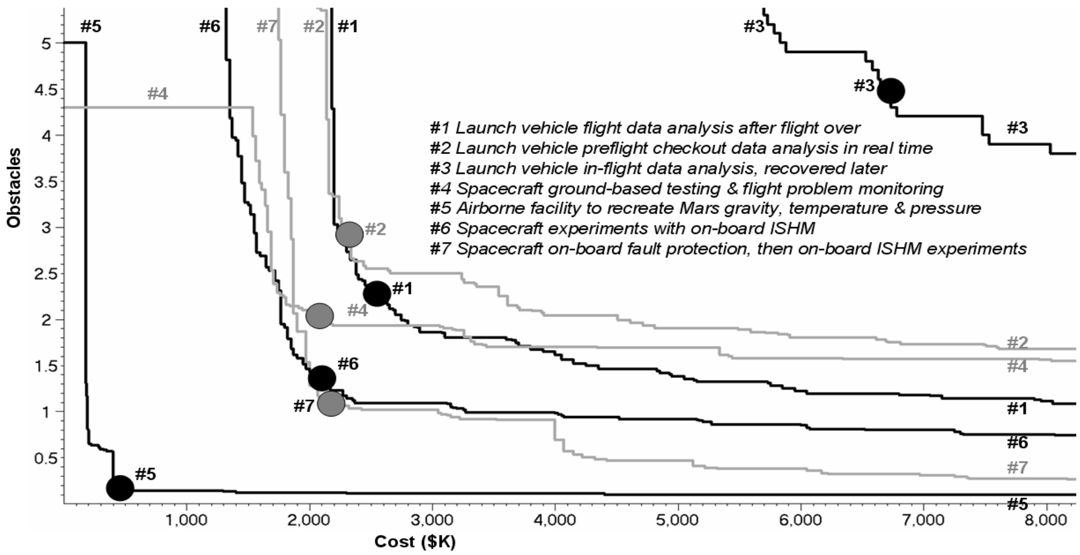

Finally, some analyses were run to look at pursuing several of the deployment opportunities at once. These showed that in most cases the combination of several opportunities' "prudent" selections continued to be near-optimal, although some instances of modest savings to be had by taking advantage of pay-once-everyone-benefit actions. Again, these results made their way into the final report.

6. Conclusions

The challenge we faced was selecting from among a number of deployment opportunities for a technology prototype, and planning for how those deployments should be planned to overcome obstacles. Lacking a standard methodology to make these kinds of decisions and plans, we applied an assessment process designed for scrutiny of a single technology application, judiciously streamlining it to apply to this study's multiplicity of opportunities. We did so cognizant of the risks such streamlining might pose, so as to retain confidence in its conclusions.

The ISHM experts used the information from this study to understand the status of the *obstacles* each opportunity would face, and to identify reasonable-cost approaches to overcoming those obstacles. This information, coupled with their own knowledge of the merits of the various opportunities, allowed them to make their final decisions of which to pursue.

7.Acknowledgments

The research described in this paper was carried out at NASA Ames Research Center and at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. In addition to the authors of this paper, the following individuals' contributions were key to the success of the study reported herein: from NASA Ames, Scott Christa, Ann Patterson-Hine, Dougal Maclise, Robert Mah, Dwight Sanderfer and Mark Schwabacher; from JPL, Arthur Murphy. We also thank Steve Cornford (JPL), who invented the methodology that underpins the TIMA process and gave us valuable guidance in this study. Steve Cornford and Martin Feather's work at JPL to develop the methodology and the custom software that supports it was funded initially by NASA Code Q (now OSMA) and subsequently by NASA ESMD.

7. References

[1] Johnson, S., "Introduction to ISHEM," *Proceedings of the First International Forum on Integrated System Health Engineering and Management in Aerospace*, Napa, California, November 2005, edited by Stephen Johnson, NASA, available from: http://ti.arc.nasa.gov/projects/ishem/ishem_definition.pdf

[2] R. Mackey, R. Some & A. Aljabri "Readiness Levels for Spacecraft Information Technologies", IEEE Aereospace Conference, 2003, pp. I-391 – I-398.

[3] Feather, M.S., Cornford, S.L., Hicks K.L. and Johnson, K.H., "Applications of tool support for risk-informed requirements reasoning," Computer Systems Science and Engineering (CRL Publishing Ltd); 20(1): January 2005 pp. 5-17.

[4] Cornford, S.L., "Managing Risk as a Resource using the Defect Detection and Prevention process," in *Proceedings of the 4th International Conference on Probabilistic Safety Assessment and Management*, 13-18 September 1998, New York City, NY, edited by A. Mosleh and R.A. Bari, International Association for Probabilistic Safety Assessment and Management.